Journal of Machine Learning Research 1 (2016) 1-33 Submitted 11/16; Published XX/XX

Specialized Support Vector Machines for Open-set Recognition

Pedro Ribeiro Mendes Júnior PEDRORMJUNIOR@GMAIL.COM
RECOD Lab., Institute of Computing (IC)
University of Campinas (UNICAMP)
Av. Albert Einstein, 1251, Campinas, SP, 13083-852, Brazil

Terrance E. Boult TBOULT@VAST.UCCS.EDU
VAST Lab., Department of Computer Science, Engineering Building
University of Colorado Colorado Springs (UCCS)
1420 Austin Bluffs Parkway, Colorado Springs, CO 80918, USA

Jacques Wainer WAINER@IC.UNICAMP.BR
Anderson Rocha ANDERSON.ROCHA@IC.UNICAMP.BR
RECOD Lab., Institute of Computing (IC)
University of Campinas (UNICAMP)
Av. Albert Einstein, 1251, Campinas, SP, 13083-852, Brazil

Editor:

arXiv:1606.03802v2 [cs.LG] 1 Nov 2016

Abstract

Often, when dealing with real-world recognition problems, we do not need, and often cannot have, knowledge of the entire set of possible classes that might appear during operational testing. Moreover, sometimes some of these classes may be ill-sampled, not sampled at all or undefined. In such cases, we need to think of robust classification methods able to deal with the "unknown" and properly reject samples belonging to classes never seen during training. Notwithstanding, almost all existing classifiers to date were mostly developed for the closed-set scenario, i.e., the classification setup in which it is assumed that all test samples belong to one of the classes with which the classifier was trained. In the open-set scenario, however, a test sample can belong to none of the known classes and the classifier must properly reject it by classifying it as unknown. In this work, we extend upon the well-known Support Vector Machines (SVM) classifier and introduce the Specialized Support Vector Machine (SSVM), which is suitable for recognition in open-set setups. The SSVM balances the empirical risk and the risk of the unknown and ensures that the region of the feature space in which a test sample would be classified as known (one of the known classes) is always bounded, ensuring a finite risk of the unknown. The same cannot be guaranteed by the traditional SVM formulation, even when using the Radial Basis Function (RBF) kernel. In this work, we also highlight the properties of the SVM classifier related to the open-set scenario, and provide necessary and sufficient conditions for an RBF SVM to have bounded open-space risk. An extensive set of experiments compares the proposed method with existing solutions in the literature for open-set recognition and the reported results show its effectiveness.

Keywords: open-set recognition, support vector machines, specialized support vector machines, bounded open-space risk, risk of the unknown

©2016 Pedro R. Mendes Júnior, Terrance E. Boult, Jacques Wainer, and Anderson Rocha.

1. Introduction

The machine learning literature is rich with works proposing classifiers for pattern recognition, with well-known examples including k-Nearest Neighbors (kNN; Bishop, 2006), Random Forests (Breiman, 2001), and the Support Vector Machines (SVM; Vapnik and Chervonenkis, 1964; Cortes and Vapnik, 1995). These classifiers were inherently designed to work in closed-set scenarios, i.e., the scenario in which all test samples must belong to a known class used in training. What happens when the test sample belongs to a class with which the classifier was not trained? Consider a scenario in which a biologist wants to classify three species of fish inside an aquarium that contains only those three species. Her aim might be counting the individuals for each species. A closed-set classifier is suitable for this task. However, consider now that the aquarium adds other fish while the biologist continues to maintaining her interest in only the original three species. The behavior of the previously designed closed-set classifier is unpredictable in this new setup. Fish belonging to an *unknown* species will be classified as belonging to one of the three known species.

We often face applications that, on operation, see examples of classes that were not seen during training. In those scenarios, a closed-set classifier would misclassify all those samples as one of the known classes. These settings require a proper open-set classifier, one that will reject such samples and classify them as unknown. We refer to such cases as open-set scenarios.

Classification in an open-set scenario, or more generally, open-set recognition, has received substantial attention lately due to its importance in real applications. In almost all the literature on pattern classification/recognition thus far, the experiments were performed considering only the closed-set scenario, i.e., all test samples belong to one of the classes with which the classifier was trained. These experiments do not well simulate most real setups. Recent works have been focused on designing new open-set algorithms and performing open-set experiments in which the classifier is tested with unseen classes (Scheirer et al., 2011, 2014; Jain et al., 2014; Bendale and Boult, 2015; Costa et al., 2014).

One possible way to address the classification in an open-set scenario is to use a closed-set classifier, obtain the similarity score to the most likely class, apply a threshold on that similarity score, and classify as unknown any test sample whose similarity score is below the threshold (Bishop, 2006). Distance rather than similarity can also be used for thresholding. However, it is well known that such thresholds are difficult to find, specially in high-dimensional spaces (Hughes, 1968).

One can also try to deal with open-set scenarios by using the SVM classifier with the one-vs-all regime to form a multiclass-from-binary approach (Rocha and Goldenstein, 2014). In the case in which all the binary classifiers that compose the multiclass one-vs-all classifier classify as negative, the test sample can be classified as unknown. Unfortunately, prior work provided no reliable policy to ensure unknowns will be classified as negative by all classifiers.

Recent works (Scheirer et al., 2013, 2014) formalized the concept of open space to indicate the region of the feature space outside the support of the training samples. Likewise, the authors in Scheirer et al. (2013, 2014) also formalized the notion of open-set risk, introducing some algorithms for open-set recognition. From those works, the concept of *positively labeled open space* (PLOS), in the binary classification level, refers to the region of the feature space outside the support of the training samples in which a test sample

would be classified as positive for the given class, i.e., the intersection of the open space and the region classified as positive. Building upon this founding block, here we refer to *known labeled open space* (KLOS) as the region of the feature space outside the support of the training samples in which a test sample would be classified as one of the known classes. The KLOS is the equivalent to PLOS for the multiclass point of view. Many classifiers proposed in the literature, e.g., kNN, Random Forests, and SVM, can generally have unbounded KLOS, and none of these can ensure the KLOS is bounded at all times. For some of them, the KLOS is the whole feature space, consequently, any test sample is always classified as one of the known classes (never as unknown).

In general, recent works (Scheirer et al., 2013, 2014; Jain et al., 2014; Costa et al., 2014) have relied upon the one-vs-all approach for multiclass-from-binary extension, aiming at decreasing the PLOS for each binary classifier and, consequently, decrease the KLOS. However, most of those methods are not able to ensure a bounded PLOS for the binary classifiers. For the one-vs-all approach, if a bounded PLOS is ensured for each binary classifier, a bounded KLOS is ensured in the multiclass level.

In this work, we introduce an extension of the traditional SVM formulation, called Specialized Support Vector Machines (SSVM), which is able to optimize the bias term (*risk of the unknown* (Scheirer et al., 2013)) and to ensure a bounded PLOS and KLOS (in the multiclass level). The main idea is to ensure a negative bias term on the SVM formulation, for each binary classifier. For SVM with the Radial Basis Function (RBF) kernel, the negative bias term will ensure that samples outside the support of the training samples shall be classified as *negative*. These samples outside the support of the training examples will be classified as negative for all binary classifiers trained with the one-vs-all approach, and therefore classified globally as unknown. We also formulate an open-set grid-search approach suitable for the open-set recognition and use it with our proposed open-set method as well as with other existing methods in the literature.

The remaining of this paper is organized as follows. In Section 2, we discuss some of the most important previous work in open-set recognition. In Section 3, we introduce the SSVM while, in Section 4, we present the experiments that validate the proposed method. Finally, in Section 5, we present the conclusions and future work.

2. Related Work

In this section, we review recent works that explicitly deal with open-set recognition in the literature. We note that other insights presented in many existing works can be somehow modified to be used for the open-set scenario. Most of those works, however, did not perform the experiments with appropriate open-set recognition setup.

In Heflin et al. (2012) and Pritsos and Stamatatos (2013), the authors present a multiclass SVM classifier based on the One-Class SVM (OCSVM; Schölkopf et al., 1999). For each of the training classes, they fit an OCSVM. In the prediction phase, all n OCSVMs classify the test sample, in which n is the number of available classes for training. The test sample is classified as the class in which its OCSVM classified as positive. When no OCSVM classifies as positive, the test sample is classified as unknown. In Heflin et al. (2012), the authors deal with multiple class classification, then when two or more OCSVMs classify as positive, the test sample is classified as belonging to those positive classes. Differently, Prit-

sos and Stamatatos (2013) choose the more confident classifier among the ones that classify as positive. In those works, the OCSVM is used with the RBF kernel. In this paper, we use the method of Pritsos and Stamatatos (2013) for comparison, hereinafter referred to as Multiclass One-Class Support Vector Machines (SVM^{OC}).

In Costa et al. (2012, 2014), the authors propose the Decision Boundary Carving (DBC), an extension of the SVM classifier aiming at a more restrictive specialization on the positive class of the binary classifier. For this, they move the hyperplane a value ϵ towards the positive class (in rare cases backwards). The value ϵ is obtained by minimizing the *training data error*. For multiclass classification, the one-vs-all approach can be used.[1] The DBC uses the RBF kernel. In this work, we use the method of Costa et al. (2012, 2014) with the one-vs-all approach for comparison. The test sample is classified as unknown when no binary classifier classifies it as positive. Complementarily, it is classified as the most confident class when one or more classifiers tags it as positive. The confidence is obtained based on the distance of the test sample to the hyperplane: the more distant, the more confident.

In Scheirer et al. (2013), the authors formalized the open-set recognition problem and proposed an extension upon the SVM classifier called 1-vs-Set Machine (1VS). Similar to the works of Costa et al. (2012, 2014), they move the main hyperplane either direction depending on the open-space risk. In addition, a second hyperplane, parallel to the main one, is created such that the positive class is between the two hyperplanes. This second hyperplane allows the samples "behind" the positive class to be classified as negative. Then a refinement step is performed on both hyperplanes to balance open-space risk and *empirical risk*. According to the authors, the method works better with the linear kernel, as the second plane does not provide much benefit for an RBF kernel which has a naturally occurring upper bound. A one-vs-all approach is used to combine the binary classifiers for open-set multiclass classification.

In Scheirer et al. (2014), the authors propose the Weibull-calibrated Support Vector Machines (WSVM). The authors proposed the Compact Abating Probability (CAP) model which decreases the probability of a test sample to be considered as belonging to one of the known classes when it is far away from the training samples. They use two stages for classification: a CAP model based on a one-class classifier followed by a binary classifier with Extreme Value Theory (EVT) normalization. The binary classifier seeks to improve discrimination and its normalization has two steps. The first aims at obtaining the probability of a test sample to belong to a positive/known class and the second step estimates the probability of it not being from the negative classes. The product of both probabilities is the probability of the test sample to belong to a positive/known class. The WSVM uses the RBF kernel and also the one-vs-all approach.

In Jain et al. (2014), the authors propose the Support Vector Machines with Probability of Inclusion (PISVM), also based on the EVT. It is an algorithm for estimating the unnormalized posterior probability of class inclusion. For each known class, a Weibull distribution (Coles, 2001) is estimated based on the smallest decision values of the positive

1. Despite dealing with a multiclass problem, Costa et al. (2012, 2014) evaluated their method in the binary fashion by obtaining the accuracy of individual binary classifiers. They did not present the multiclass version of the classifier directly. Therefore, in this work, we consider their method with the one-vs-all approach in the experiments.

training samples. The binary classifier for each class is an SVM with RBF kernel trained using the one-vs-all approach, i.e., the samples of all remaining classes are considered as negative samples. They introduce the idea of *cross-class validation* which is similar to our open-set grid search, which we shall discuss in details in Section 3.4. For a test sample, PISVM chooses the class for which the decision value produces the maximum probability of inclusion. If that maximum is below a given threshold, the input is marked as unknown.

3. Specialized Support Vector Machines (SSVMs)

One cannot ensure that the positive class of the traditional SVM has a bounded PLOS, even when the RBF kernel is used. The main characteristic of the proposed SSVM is that a high enough regularization parameter ensures a bounded PLOS for every known class of interest. The regularization parameter is a weight for optimizing the risk of the unknown in relation to the empirical risk. In Section 3.2, we present how to ensure a PLOS using the RBF kernel. In Section 3.3, we present the formulation of the optimization problem of the SSVM. In Section 3.4, we present the process of optimization of the parameters. To begin with, in Section 3.1, we present some basic aspects of the SVM classifier.

3.1 Basic aspects of Support Vector Machines

The SVM is a binary classifier that, given a set X of training samples $\mathbf{x}_i \in \mathbb{R}^d$ and the corresponding labels $y_i \in \{-1, 1\}$, $i = 1, \ldots, m$, it finds a maximum-margin hyperplane that separates $\mathbf{x}_i$ for which $y_i = -1$ from $\mathbf{x}_i$ for which $y_i = 1$ (Cortes and Vapnik, 1995). We consider the soft margin case with parameter C.

The primal optimization problem is usually defined as

$$\min_{\mathbf{w},b,\xi} \frac{1}{2}||\mathbf{w}||^2 + C\sum_{i=1}^{m}\xi_i,$$

$$\text{s.t. } y_i(\mathbf{w}^T\mathbf{x}_i + b) \geq 1 - \xi_i, \ \forall i, \tag{1}$$

$$\xi_i \geq 0, \ \forall i. \tag{2}$$

To solve this optimization problem, we use the Lagrangian method to create the dual optimization problem. In this case, the final Lagrangian is defined as

$$\mathcal{L}(\mathbf{w}, b, \xi, \alpha, r) = \sum_{i=1}^{m}\alpha_i - \frac{1}{2}||\mathbf{w}||^2, \tag{3}$$

in which $\alpha_i \in \mathbb{R}$, $r_i \in \mathbb{R}$, $i = 1, \ldots, m$, are the Lagrangian multipliers. Then, the optimization problem now is defined as

$$\min_{\alpha} W(\alpha) = -\mathcal{L}(\mathbf{w}, b, \xi, \alpha, r) = \frac{1}{2}||\mathbf{w}||^2 - \sum_{i=1}^{m}\alpha_i, \tag{4}$$

$$\text{s.t. } 0 < \alpha_i < C, \ \forall i, \tag{5}$$

$$\sum_{i=1}^{m}\alpha_i y_i = 0. \tag{6}$$

The decision function of a test sample $\mathbf{x}$ comes from the constraint in Equation (1) and is defined as

$$f(\mathbf{x}) = \text{sign}(\mathbf{w}^T\mathbf{x} + b) = \text{sign}\left(\sum_{i=1}^{m} y_i\alpha_i\mathbf{x}_i^T\mathbf{x} + b\right).$$

In Boser et al. (1992), the authors proposed a modification in SVM for the cases in which the training data are not linearly separated in the feature space. Instead of linearly separating the samples in the original space $\mathcal{X}$ of the training samples in X, the samples are projected onto a higher dimensional space $\mathcal{Z}$ in which they are linearly separated. This projection is accomplished using the kernel trick. One advantage of this method is that in addition to separating non-linear data, the optimization problem of the SVM remains almost the same: instead of calculating the inner product $\mathbf{x}^T\mathbf{x}'$, it uses a kernel $K(\mathbf{x}, \mathbf{x}')$ that is equivalent to the inner product $\phi(\mathbf{x})^T\phi(\mathbf{x}')$ in a higher dimensional space $\mathcal{Z}$, in which $\phi : \mathcal{X} \mapsto \mathcal{Z}$ is a projection function. When using the kernel trick, we do not need to know the $\mathcal{Z}$ space explicitly.

Using kernels, the decision function of a test sample $\mathbf{x}$ becomes

$$f(\mathbf{x}) = \text{sign}\left(\sum_{i=1}^{m} y_i\alpha_i K(\mathbf{x}_i, \mathbf{x}) + b\right). \tag{7}$$

The most used kernel for SVM is the RBF kernel (Schölkopf and Smola, 2001), defined as follows.

$$K(\mathbf{x}, \mathbf{x}') = e^{-\gamma||\mathbf{x}-\mathbf{x}'||^2}. \tag{8}$$

It is proved that using this kernel, the projection space $\mathcal{Z}$ is an ∞-dimensional space.

3.2 Ensuring a bounded PLOS

By simply using the RBF kernel we cannot ensure the PLOS is bounded.

Theorem 1 *SVM with RBF kernel has a bounded* positively labeled open space *(PLOS) if and only if the bias term b is negative.*[2]

Proof *We know that*

$$\lim_{d\to\infty} e^{-\gamma d^2} = 0$$

for positive values of γ. It follows from Equation (8) *that when samples $\mathbf{x}$ and $\mathbf{x}'$ are far apart, $K(\mathbf{x}, \mathbf{x}')$ tends to 0. For the cases in which the test sample $\mathbf{x}$ is far away from every support vector $\mathbf{x}_i$,*

$$\sum_{i=1}^{m} y_i\alpha_i K(\mathbf{x}_i, \mathbf{x})$$

also tends to 0. From Equation (7) *it follows that*

$$f(\mathbf{x}) \to \text{sign}\,(b)$$

2. Note in some implementations, including the LibSVM library (Chang and Lin, 2011), the decision function is defined as $f(\mathbf{x}) = \text{sign}(\mathbf{w}^T\mathbf{x} - \rho)$. In that case, instead of ensuring a negative bias term b, one must ensure a positive bias term ρ to bound the PLOS.

when $\mathbf{x}$ *is far away from the support vectors. Therefore, for negative values of* b*,* $f(\mathbf{x})$ *is always negative for far away* $\mathbf{x}$ *samples. That is, samples in an bounded region of the feature space will be classified as positive. For the only if direction, let* b *be positive. Then there will exist a distance* d *such that* $\forall i : \|\mathbf{x}_i - \mathbf{x}\| > d \implies f(\mathbf{x}) = \text{sign}(b) > 0$*, i.e., all positively classified samples will be in an unbounded region of the feature space.* ■

Figure 1 depicts the rationale behind Theorem 1. The z axis represents the decision values that possible 2-dimensional test samples (x, y) would have for different regions of the feature space. The training samples are normalized between 0 and 1. Note in the subfigures that for possible test samples far away from the training ones, e.g., $(2, 2)$, the decision value is equal to the bias term b. Note in Figure 1c that an unbounded region of the feature space would have samples classified as positive. Consequently, all those samples would be classified as class 3 by the final multiclass-from-binary classifier. In general SVM usage, both positive and negatives biases occur as b depends on the training data.

In case of the SVMs without explicit bias term (Vogt, 2002; Kecman et al., 2005), $b = 0$ is implicit.[3] Consequently, the decision function is defined as

$$f(\mathbf{x}) = \text{sign}\left(\sum_{i=1}^{m} y_i \alpha_i K(\mathbf{x}_i, \mathbf{x})\right).$$

For test samples far away from the support vectors, we have that $\sum_{i=1}^{m} y_i \alpha_i K(\mathbf{x}_i, \mathbf{x})$ converges to 0 from the bottom or from the above, depending on the training samples. Consequently, a bounded PLOS cannot be ensured in all cases.

Theorem 1 also provides a solution to the problem of unbounded PLOS. We can ensure a bounded PLOS by using an RBF kernel and ensuring a negative b. In Section 3.3, we present a new SVM optimization objective that optimizes the margin while ensuring the bias term b is negative.

3.3 Optimization to ensure negative bias term b

As we discussed in Section 3.2, we must ensure a negative b to obtain a bounded PLOS. For this, we define the SSVM optimization problem as

$$\min_{\mathbf{w},b,\xi} \frac{1}{2}||\mathbf{w}||^2 + C\sum_{i=1}^{m} \xi_i + \lambda b, \tag{9}$$

subject to the same constraints defined in Equations (1) and (2), in which λ is a regularization parameter that trades off between the empirical risk and the risk of the unknown.

From Equation (9), the dual formulation has the same Lagrangian defined in Equation (3). Consequently, we have to optimize the same function as defined in Equation (4) with the constraint in Equation (5). However, the constraint in Equation (6) is replaced by the constraint

$$\sum_{i=1}^{m} \alpha_i y_i = \lambda. \tag{10}$$

3. The main difference from the SVM without bias term to the traditional SVM is that the constraint in Equation (6) does not exist in the dual formulation.

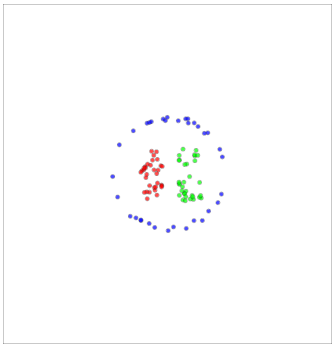

Boat data set with 3 classes: red (the central class to the left), green (the central class to the right), and blue (the class with the ring shape).

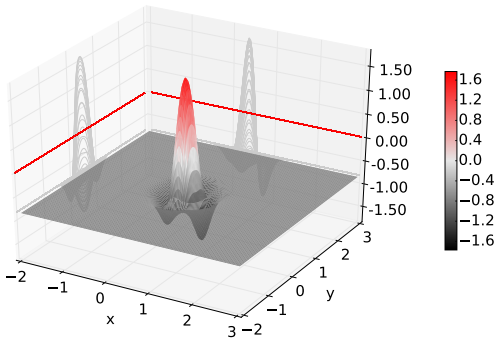

(a) Class 1 (red). $b = -0.832$.

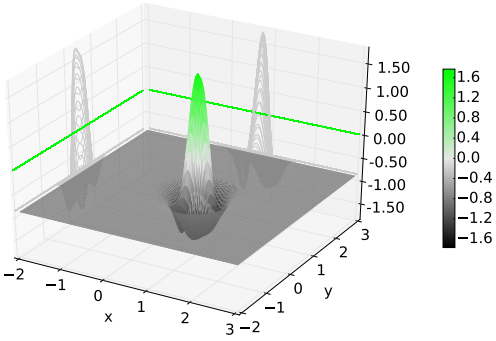

(b) Class 2 (green). $b = -0.86$.

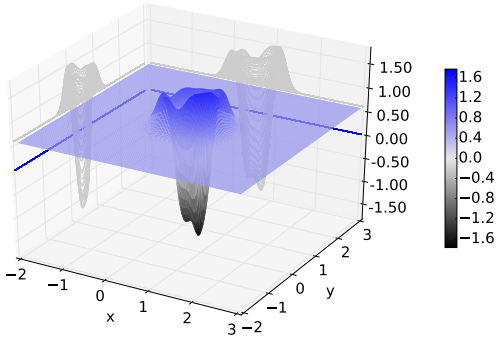

(c) Class 3 (blue). $b = +0.594$.

Figure 1: Behavior of the Support Vector Machines (SVM) with a Radial Basis Function (RBF) kernel on a 2-dimensional synthetic data set. The image on the top-left depicts the Boat data set. Figures (a)–(c) correspond to the red, green and blue classes of the Boat data set, respectively. The x, y axis in Figures (a)–(c) represent the two features of the data set (the data set is normalized between 0 and 1). The z axis shows the value of the decision function $\sum_{i=1}^{m} y_i \alpha_i K(\mathbf{x}_i, \mathbf{x}) w + b$ (Equation 7 without sign function) and the colored lines in the walls depict the point 0, that separates the positive class from the negative one (equivalent to the sign function of Equation 7). Note in Figure (c) that an unbounded region of the feature space remains in the positive side, as $b > 0$ and $f(x, y) \approx b$ for far away (x, y) points.

The same Sequential Minimal Optimization (SMO) algorithm proposed by Platt (1998), with the Working Set Selection (WSS) proposed by Fan et al. (2005), for optimizing ensuring the constraint in Equation (6) can be applied to this optimization containing the constraint of the Equation (10). As the main idea of the SMO algorithm is to ensure that $\sum \alpha_i y_i$ remains the same from one iteration to the other, before the optimization starts, we initialize α_i such that $\sum \alpha_i y_i = \lambda$. For this, we let $\alpha_i = \lambda/m_p$, $\forall i$ such that $y_i = 1$, where m_p is the number of positive training samples.

Proposition 2 *In the soft margin case, with parameter C, the maximum valid value for λ is Cm_p.*
Proof *From Equation (5), $0 \leq \alpha_i \leq C$. The maximum value $\lambda = \sum \alpha_i y_i$ is thus obtained by setting $\alpha_i = C$ for i such that $y_i = 1$ and setting $\alpha_i = 0$ for i such that $y_i = -1$. This yields $\lambda \leq Cm_p$* ■

Thus during optimization, we must ensure $\lambda \leq Cm_p$, otherwise, for $\lambda > Cm_p$, the constraint in Equation (5) would be broken for some α_i.

Despite Proposition 2 saying that it is allowed $\lambda = Cm_p$, when it happens, we have that $\alpha_i = C$ for $y_i = 1$ and $\alpha_i = 0$ for $y_i = -1$, and there will be no optimization. In this case, despite satisfying the constraints, there is no flexibility for changing the values of α_i because, for each pair α_i, α_j selected by the WSS algorithm (Fan et al., 2005), we must update $\alpha_i = \alpha_i + \nabla_\alpha$, $\alpha_j = \alpha_j + \nabla_\alpha$ when $y_i \neq y_j$ and $\alpha_i = \alpha_i - \nabla_\alpha$, $\alpha_j = \alpha_j + \nabla_\alpha$ when $y_i = y_j$. For any $\nabla_\alpha \neq 0$, the constraint $0 < \alpha_i < C$ would break for either α_i or α_j, for any pair selected. Then, in practice, we grid search λ in the interval $0 \leq \lambda < Cm_p$.

Proposition 3 *There exists some λ such that we can obtain a bias term $b < 0$.*
Proof *From the KarushKuhnTucker (KKT; Bishop, 2006) conditions, the bias term is defined as*

$$\begin{aligned} b &= y_i - \sum_{j=1}^{m} \alpha_j y_j K(\mathbf{x}_i, \mathbf{x}_j) \\ &= y_i - \sum_{\substack{j=1: \\ y_j=1}}^{m} \alpha_j K(\mathbf{x}_i, \mathbf{x}_j) + \sum_{\substack{j=1: \\ y_j=-1}}^{m} \alpha_j K(\mathbf{x}_i, \mathbf{x}_j), \end{aligned}$$

*for any i such that $0 < \alpha_i < C$. Now, let's consider two possible cases: (1) $y_i = 1$ and (2) $y_i = -1$. For **Case (1)**, we have*

$$b = 1 - \alpha_i - \sum_{\substack{j=1: \\ y_j=1, \\ j\neq i}}^{m} \alpha_j K(\mathbf{x}_i, \mathbf{x}_j) + \sum_{\substack{j=1: \\ y_j=-1}}^{m} \alpha_j K(\mathbf{x}_i, \mathbf{x}_j),$$

as $K(\mathbf{x}_i, \mathbf{x}_i) = 1$. Note that $0 < K(\mathbf{x}, \mathbf{x}') \leq 1$ for the RBF kernel. To show that there exists a λ such that $b < 0$, we analyze the worst case, i.e., when the kernel in the second

summation is 1*. Then, we have*

$$b = 1 - \alpha_i - \sum_{\substack{j=1: \\ y_j=1, \\ j\neq i}}^{m} \alpha_j K\left(\mathbf{x}_i, \mathbf{x}_j\right) + \sum_{\substack{j=1: \\ y_j=-1}}^{m} \alpha_j.$$

According to Equation (10)*, we have*

$$\sum_{\substack{j=1: \\ y_j=-1}}^{m} \alpha_j = \sum_{\substack{j=1: \\ y_j=1}}^{m} \alpha_j - \lambda, \tag{11}$$

then

$$b = 1 - \sum_{\substack{j=1: \\ y_j=1, \\ j\neq i}}^{m} \alpha_j K\left(\mathbf{x}_i, \mathbf{x}_j\right) + \sum_{\substack{j=1: \\ y_j=1 \\ j\neq i}}^{m} \alpha_j - \lambda$$

Analyzing the worst case again, considering the $\alpha_j = C$*, we have*

$$\begin{aligned} b &= 1 - C \sum_{\substack{j=1: \\ y_j=1, \\ j\neq i}}^{m} K\left(\mathbf{x}_i, \mathbf{x}_j\right) + C\left(m_p - 1\right) - \lambda \\ &= 1 + Cm_p - C - C \sum_{\substack{j=1: \\ y_j=1, \\ j\neq i}}^{m} K\left(\mathbf{x}_i, \mathbf{x}_j\right) - \lambda \end{aligned}$$

Then, to ensure $b < 0$ *it is sufficient to let*

$$\lambda > 1 + Cm_p - C \left(1 + \sum_{\substack{j=1: \\ y_j=1, \\ j\neq i}}^{m} K\left(\mathbf{x}_i, \mathbf{x}_j\right) \right).$$

Given a $C \geq 1$ *it is always possible to obtain a* λ *such that* $\lambda < Cm_p$*. In fact, this bound is not tight and it is only necessary to have a* C *such that the absolute value of the last term in the formula is greater than* 1*, and often* $C < 1$ *still supports a viable* λ*.*

For ***Case (2)****, we have*

$$b = -1 - \sum_{\substack{j=1: \\ y_j=1}}^{m} \alpha_j K\left(\mathbf{x}_i, \mathbf{x}_j\right) + \sum_{\substack{j=1: \\ y_j=-1}}^{m} \alpha_j K\left(\mathbf{x}_i, \mathbf{x}_j\right).$$

Again, considering the worst case for the values of the kernel and using the equality in Equation (11)*, we have*

$$b = -1 - \sum_{\substack{j=1: \\ y_j=1}}^{m} \alpha_j K\left(\mathbf{x}_i, \mathbf{x}_j\right) + \sum_{\substack{j=1: \\ y_j=1}}^{m} \alpha_j - \lambda.$$

Again, considering the greatest possible value for b, by considering $\alpha_j = C$, we have

$$b = -1 - C \sum_{\substack{j=1: \\ y_j=1}}^{m} K(\mathbf{x}_i, \mathbf{x}_j) + Cm_p - \lambda.$$

In this case, to ensure $b < 0$ it is sufficient to let

$$\lambda > Cm_p - 1 - C \sum_{\substack{j=1: \\ y_j=1}}^{m} K(\mathbf{x}_i, \mathbf{x}_j),$$

which is possible to obtain for any value of C. ■

Note that in Proposition 3 we considered a very extreme case for the proof. In Case (1), for example, considering $y_i = 1$, we considered $K(\mathbf{x}_i, \mathbf{x}_j) = 1$ for $y_j = -1$ and $K(\mathbf{x}_i, \mathbf{x}_j) \approx 0$ for $y_j = 1$. It means that all the negative samples have the same feature vector of sample $\mathbf{x}_i$ under consideration and all positive samples are far away from sample $\mathbf{x}_i$. In practice, we do not have the λ nearly as constrained as in the proof to ensure a negative bias term. More than this, in our experiments with the SVM, we observed that most of the time the bias term is negative for a binary classifier trained with the one-vs-all approach, i.e., it is often the case that even with $\lambda = 0$ the bias will be negative. More details about this behavior is shown in Section 4.5.

In Appendix A we present the complete formulation of the optimization problem for the SSVM classifier.[4]

Choosing the λ parameter for the SSVM The Proposition 3 states that we can find a λ parameter that ensures a bounded PLOS for the optimization problem presented above. To ensure this, models with a non-negative bias term receive accuracy of $-\infty$ on the validation set, during the grid search. Nevertheless, it can happen that certain λ allows a negative bias term during the grid search but not during training in a larger set of training samples. In this case, once the parameters are obtained by grid search, if the obtained λ does not ensure a negative bias term for training set, we can retrain the classifier with an increased value for λ until a negative bias term is obtained in the model. However, in our experiments it was not necessary. The values of λ obtained during grid search were able to ensure a negative bias term for all binary classifiers.

3.4 Grid search in an open-set setup

Two different approaches for grid searching the parameters in multiclass-from-binary classifiers have been considered in the literature thus far (Chang and Lin, 2011): (1) the external approach and (2) the internal one. As these forms were defined in a closed-set setting, we properly extend them to open-set scenarios. In the external approach, we grid search the

4. The source code, extended upon the LibSVM implementation (Chang and Lin, 2011), will be available upon acceptance of this paper.

parameters in the multiclass level so that all binary classifiers that compose the multiclass classifiers share the same set of parameters. On the other hand, in the internal approach, each binary classifier performs its own grid search (Chang and Lin, 2011).

To perform a grid search on the training set X, a subset $F \subset X$ is used to fit an SVM classifier while performing the grid search. The subset $V = X - F$ then forms the validation set. This can be done for both external and internal approaches.

There are two forms of dividing the training set X into the fitting set F and the validation set V while performing the grid search. In this section, we present what we call the closed-set form (the traditional form) and the open-set form. Both forms can be applied to both external and internal approaches.

As we consider the open-set setup in our experiments, the classes represented in the set X are a proper subset of the classes of the data set. It means that some unknown classes appear during testing. For both the closed-set and the open-set grid search approaches, we consider open-set scenarios in the testing phase. The difference of the open-set grid search is that it better simulates an open-set setup while performing the grid search during training.

External closed-set grid search As an example of the external closed-set grid search, for the multiclass classifier, 80% of the samples in X are randomly selected for F and 20% for V. F can have representative samples of all n available classes. Then, the multiclass classifier is trained with F and evaluated with V, testing possible parameters. This is the traditional form of the external grid search: we do not mind when the negative samples in F represent the n classes available for training.

External open-set grid search In the external open-set grid search, for the multiclass classifier, $\lfloor n/2 \rfloor$ of the n available classes for training are taken for V. Among the samples of the remaining $\lceil n/2 \rceil$ classes, 80% are randomly selected for F and 20% for V. Then, the multiclass classifier is trained with F and evaluated with V, testing possible parameters.

Internal closed-set grid search In the internal closed-set grid search, for a binary classifier, 80% of the samples in X are randomly selected for F and 20% for V. We ensure a random partitioning such that both F and V have representative samples of the positive and negative classes. In this case, the negative samples in V can have representative samples of all $n-1$ negative classes, as the partitioning between F and V is random, i.e., when using V to obtain the accuracy during the grid search, all classes in V can be known a priori by the classifier trained with F. Then, the binary classifier is trained with F and evaluated with V, testing possible parameters. This is the traditional form of the internal grid search: we do not mind when the negative samples in F represent the $n-1$ negative classes of the one-vs-all training.

Internal open-set grid search In the internal open-set grid search, for a binary classifier, $\lfloor (n-1)/2 \rfloor$ of the $n-1$ negative classes are taken for V. Among the samples of the remaining $\lceil (n-1)/2 \rceil + 1$ classes (considering the positive class), 80% are randomly selected for F and 20% for V. Also, a random partitioning such that both F and V have representative samples of the positive and negative classes must be guaranteed. Then, the binary classifier is trained with F and evaluated with V, testing possible parameters.

Figure 2 depicts how the set X of training samples is partitioned for the closed- and open-set grid search.

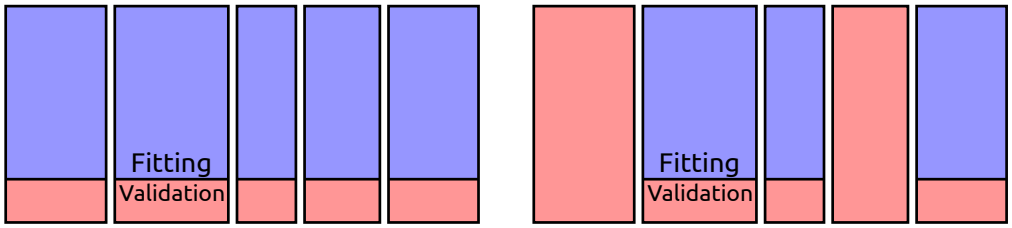

(a) Closed-set grid search. (b) Open-set grid search.

Figure 2: Difference between closed- and open-set grid search. Each colored box represents a known class during training. The blue color represents the amount of samples used to train a classifier during the grid-search procedure (fitting set). The red part represents the validation set used to evaluate the classifier. The parameters with the best accuracy on the validation set are selected for training the classifier. For the open-set grid search depicted in (b), some of the known classes are considered as unknown during the grid-search procedure. For both external and internal grid search, the closed- and open-set variations are split the same way, however ensuring, for the case of the internal grid search, the positive class have representative samples for the fitting set.

4. Experiments

In this section, we present the experiments and details for comparing the proposed method with the existing ones in the literature, discussed in Section 2. For a fair comparison, we also implemented the closed- and open-set grid search described in Section 3.4 for the existing methods in the literature considered for comparison in this paper. We denote a given classifier C using the closed- and open-set grid search as C_C and C_O, respectively.

All compared methods are SVM-based and use the one-vs-all approach in the multiclass level. For most of the classifiers, we used the internal grid search approach. For the classifiers in which we have the multiclass implementation provided by the authors (e.g., 1VS, WSVM, and PISVM), we use the external approach for grid-searching.

4.1 Evaluation measures

Most of the proposed evaluation performance measures in the literature are focused on binary classification, e.g., traditional classification accuracy, f-measure, etc. (Sokolova and Lapalme, 2009). As a matter of fact, some measures can be easily adapted to multiclass classification (e.g., average accuracy, multiclass f-measure, etc.; Sokolova and Lapalme, 2009), however, there is no specific measure for the open-set recognition problem to date. Note that using these measures, e.g., the multiclass f-measure, is not appropriate, as the *unknown* is not a single class and the amount of unknown samples can be much larger than the samples of the known classes. In this section, we use an extension of the average accuracy and multiclass f-measure for evaluating experiments in open-set setups. With these new measures we aim at gathering a better view of the classification in open-set scenarios.

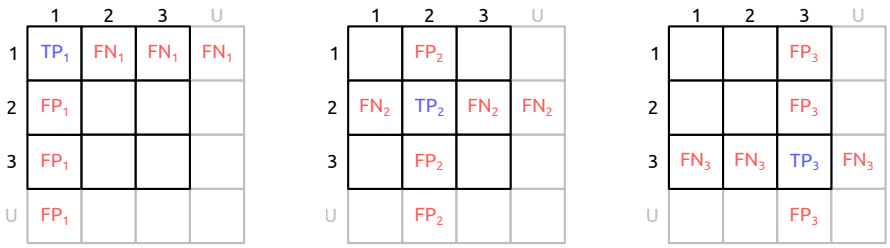

(a) (b) (c)

Figure 3: Confusion matrix used to calculate the open-set f-measure (OSFM). Example related to the computation of the precision and recall measures from the confusion matrix with three known classes (1, 2, and 3) and the unknown (U). The FN_i in the column U and FP_i in the row U account for *false unknown* and *false known*, respectively. The cell in the intersection of column U and row U is *not* regarded as true positive in the same way it would be considered by the multiclass closed-set f-measure.

In this vein, the accuracy on known samples (AKS) is the accuracy obtained on the testing samples that belong to one of the classes with which the classifier was trained. The accuracy on unknown samples (AUS) is the accuracy on the testing samples whose classes have no representative samples in the training set. The normalized accuracy (NA) treats known and unknown equally and is computed from the AKS and the AUS as

$$\mathrm{NA} = \frac{\mathrm{AKS} + \mathrm{AUS}}{2}.$$

The open-set f-measure (OSFM), despite considering the multiclass classification problem, takes into account the classification of the unknown samples. For this, an additional line and column referring to the unknown is added to the confusion matrix used to calculate the OSFM, as Figure 3 depicts.

Then the formulation of the OSFM is the same of the traditional f-measure,

$$\mathrm{OSFM} = \frac{2 \times \mathrm{precision} \times \mathrm{recall}}{\mathrm{precision} + \mathrm{recall}}.$$

The difference is in the definition of the precision and recall. See Equations (12) and (13).

$$\begin{aligned} \mathrm{precision}_M &= \frac{\sum_{i=1}^{l-1} \frac{\mathrm{TP}_i}{\mathrm{TP}_i + \mathrm{FP}_i}}{l-1}, \\ \mathrm{recall}_M &= \frac{\sum_{i=1}^{l-1} \frac{\mathrm{TP}_i}{\mathrm{TP}_i + \mathrm{FN}_i}}{l-1}, \end{aligned} \tag{12}$$

$$\begin{aligned} \mathrm{precision}_\mu &= \frac{\sum_{i=1}^{l-1} \mathrm{TP}_i}{\sum_{i=1}^{l-1} (\mathrm{TP}_i + \mathrm{FP}_i)}, \\ \mathrm{recall}_\mu &= \frac{\sum_{i=1}^{l-1} \mathrm{TP}_i}{\sum_{i=1}^{l-1} (\mathrm{TP}_i + \mathrm{FN}_i)}, \end{aligned} \tag{13}$$

in which $l = n + 1$ is the number of rows and columns of the confusion matrix (represented in Figure 3) and n is the number of classes available in the training set. TP, FP, and FN stand for true positive, false positive, and false negative, respectively. The last row and column of the confusion matrix refers to the unknown samples/classifications.

Using the precision and recall of the Equation (12), we define the macro-averaging open-set f-measure (OSFM_M) and using Equation (13) we define the micro-averaging open-set f-measure (OSFM_μ). The difference between Equations (12), (13) and the traditional definitions of precision and recall resides on the looping through the confusion matrix: not all rows and columns are covered (the row and column regarding the unknown is not covered). See Figure 3 for an example. In the same way that the traditional f-measure is invariant to the true negative (Sokolova and Lapalme, 2009), the OSFM is invariant to the true unknown. Note also that the OSFM takes into account both the false unknown (known samples classified as unknown) and the false known (unknown samples classified as known).

For a fair comparison with previous methods in the literature (Scheirer et al., 2013, 2014; Jain et al., 2014) which only showed performance figures using the traditional f-measure, in Section 4.4, we also present the results regarding the traditional multiclass f-measure (Sokolova and Lapalme, 2009) considering both macro-averaging f-measure (FM_M) and micro-averaging f-measure (FM_μ).

4.2 Data sets

For validating the proposed method and comparing it with existing methods, we consider five data sets. In the 15-Scenes (Lazebnik et al., 2006) data set (with 15 classes), the images are represented by a bag-of-visual-word vector created with soft assignment (van Gemert et al., 2010) and max pooling (Boureau et al., 2010), based on a codebook of 1000 Scale Invariant Feature Transform (SIFT) codewords (Lowe, 2004). The Letter (Frey and Slate, 1991; Michie et al., 1994) data set (with 26 classes) represents the letters of the English alphabet (black-and-white rectangular pixel displays). In the Auslan (Kadous, 2002) data set (with 95 classes), the data was acquired using two Fifth Dimension Technologies (5DT) gloves hardware and two Ascension Flock-of-Birds magnetic position trackers. In the Caltech-256 (Griffin et al., 2007) data set (with 256 classes), the feature vectors consider a bag-of-visual-words characterization approach, with features acquired with dense sampling, SIFT descriptor for the points of interest, hard assignment (van Gemert et al., 2010), and average pooling (Boureau et al., 2010). Finally, for the ALOI (Geusebroek et al., 2005) data set (with 1000 classes), the features were extracted with the Border/Interior (BIC) descriptor (Stehling et al., 2002). These data sets or other data sets could be used with different characterizations, however, in this work, we focus on the learning part of the problem rather than on the feature characterization one.

In Table 1, we summarize the main features of the considered data sets in terms of number of samples, number of classes, dimensionality, and approximate number of samples per class.[5]

5. The feature vectors for the data sets used in this work will be available upon acceptance of this paper.

Data set	# samples	# classes	# features	# samples/class
15-Scenes	4485	15	1000	299
Letter	20,000	26	16	769
Auslan	146,949	95	22	1546
Caltech-256	29,780	256	1000	116
ALOI	108,000	1000	128	108

Table 1: General characteristics of the data sets used in the experiments.

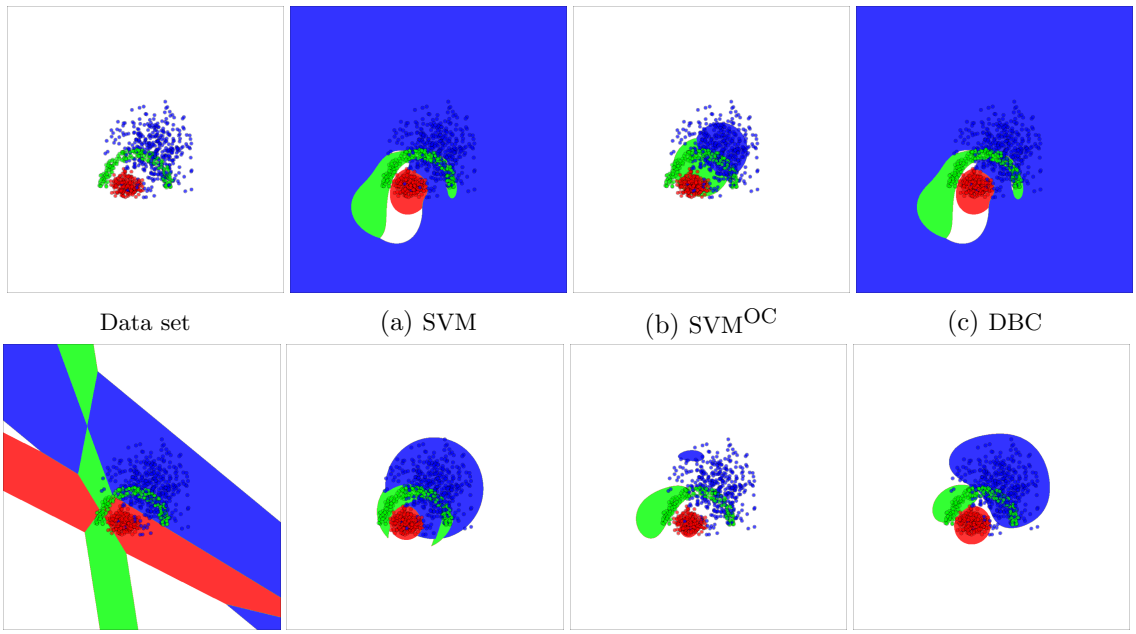

Data set (a) SVM (b) SVM^{OC} (c) DBC

(d) 1VS (e) WSVM (f) PISVM (g) SSVM

Figure 4: Decision boundaries for the Cone-torus data set. The non-white regions represent the region in which a test sample would be classified as belonging to the same class of the samples with the same color. All samples in the white regions would be classified as unknown.

4.3 Decision regions

We start with an artificial 2-dimensional data set so that the decision regions of each classifier can be visualized. For these cases, we show the region of the feature space in which a possible test sample would be classified as one of the known classes or unknown. We also show how each classifier handles the open space. Figure 4 depicts the images of *decision regions* for the Cone-torus (Kuncheva and Hadjitodorov, 2004) data set. In Figure 4g, as expected, we see that the SSVM gracefully bounds the open space; any sample that would appear in the white region would be classified as unknown.[6]

6. We present images of decision regions for other 2-dimensional data sets in Appendix B.

Measure	SVM	SVM^{OC}	DBC	1VS	WSVM	PISVM
NA	**0.0017***	**0.0000***	**0.0000***	**0.0000***	**0.0000***	**0.0131**
$OSFM_M$	**0.0003***	**0.0000***	**0.0000***	**0.0000***	**0.0000***	**0.0400**
$OSFM_\mu$	**0.0000***	**0.0000***	**0.0000***	**0.0000***	**0.0000***	**0.0023***
FM_M	**0.0006***	**0.0000***	**0.0000***	**0.0000***	**0.0005***	0.3580
FM_μ	**0.0000***	**0.0000***	**0.0000***	**0.0000***	**0.0000***	**0.0000***

Table 2: Binomial statistical tests comparing the $SSVM_O$ with other methods using the open-set grid search. Each cell in table compares the results for all data sets considering all number of available classes. **Bold** means there is statistical difference with 95% of confidence. The "*" indicates the statistical difference is with 99% of confidence.

4.4 Results

We performed a series of experiments simulating an open-set scenario in which 3, 6, 9, and 12 classes are available for training the classifiers. The remaining classes of each data set are unknown at the training phase and only appear on the testing stage. Since different data sets have a different number of known classes, the fraction of unknown classes, or the openness (in the sense of Scheirer et al., 2013), varies per data set. For each data set, method and number n of available classes, we ran 10 experiments by choosing n random classes for training, among the classes of the data set. The selected sets of classes are used across each experiment with different classifiers. In addition, the same samples used for training one classifier C_i is used when training another classifier C_j (a similar setup is adopted for the testing, and validation sets), which is referred to as a paired experiment.

Figures 5 and 6 show the results considering NA and $OSFM_M$, respectively, for all data sets.[7] We see that, in general, $SSVM_O$ maintains a good accuracy among the data sets, while the other methods have good accuracy for one or two data sets but do not maintain a similar accuracy for the others. The superiority of $SSVM_O$ is confirmed by the Binomial test we performed using the Holm method to adjust each of the paired p-values to account for multiple comparisons (Demšar, 2006). Considering all data sets and the experiments performed with 3, 6, 9, and 12 available classes together, comparing the methods with open- and closed-set grid search separately, both $SSVM_O$ and $SSVM_C$ were statistically significantly better in virtually all cases when compared to all other baseline classifiers, considering NA, $OSFM_M$, $OSFM_\mu$, FM_M, and FM_μ measures.

In Table 2, we present the p-values for comparing the $SSVM_O$ with the baselines implemented with open-set grid search. For virtually all those cases, the SSVM outperforms its counterparts with statistical significance. Only when compared to $PISVM_O$, with respect to the FM_M, is there not a statistically significant difference. When comparing the implementations with closed-set grid search, the $SSVM_C$ also obtains better results in all cases. For this comparison, the difference is significant with 95% in all cases.

7. We present additional results considering the $OSFM_\mu$, FM_M, and FM_μ in Appendix B.

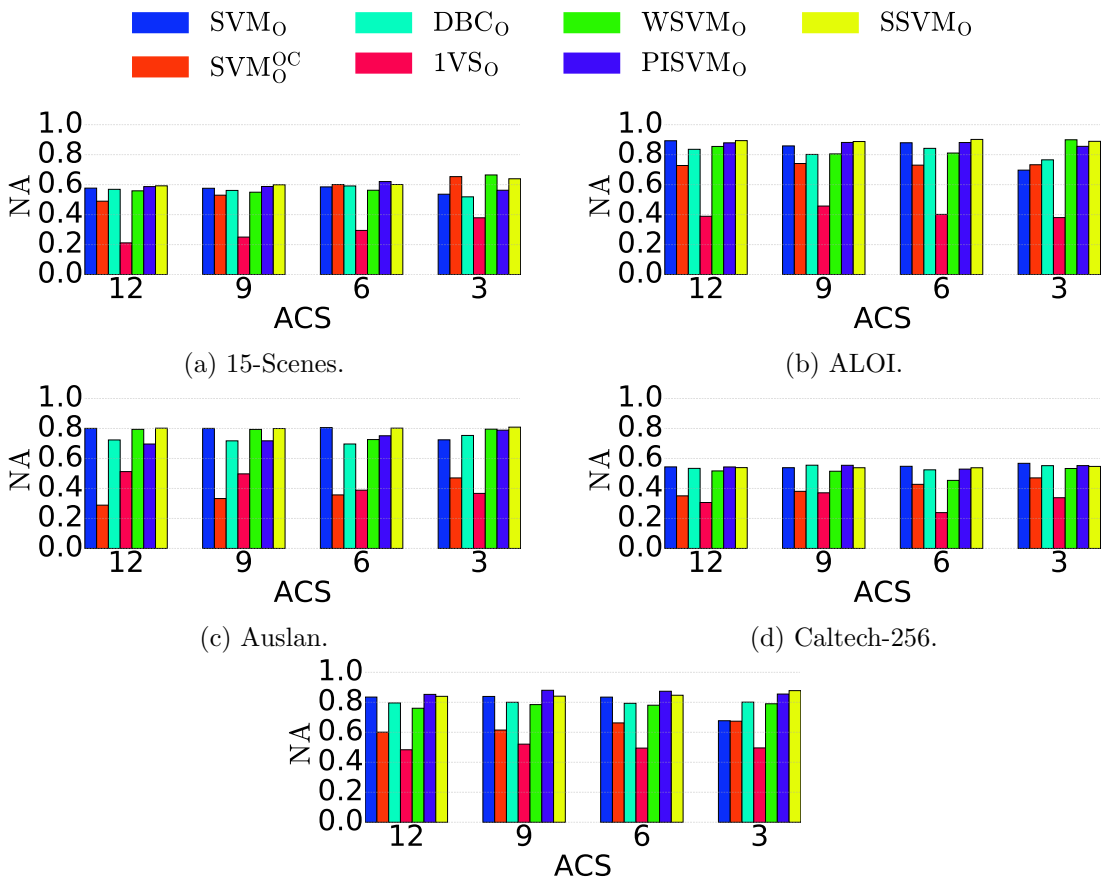

(a) 15-Scenes.

(b) ALOI.

(c) Auslan.

(d) Caltech-256.

(e) Letter.

Figure 5: Results for all data sets regarding normalized accuracy (NA) considering 3, 6, 9, and 12 available classes (ACS).

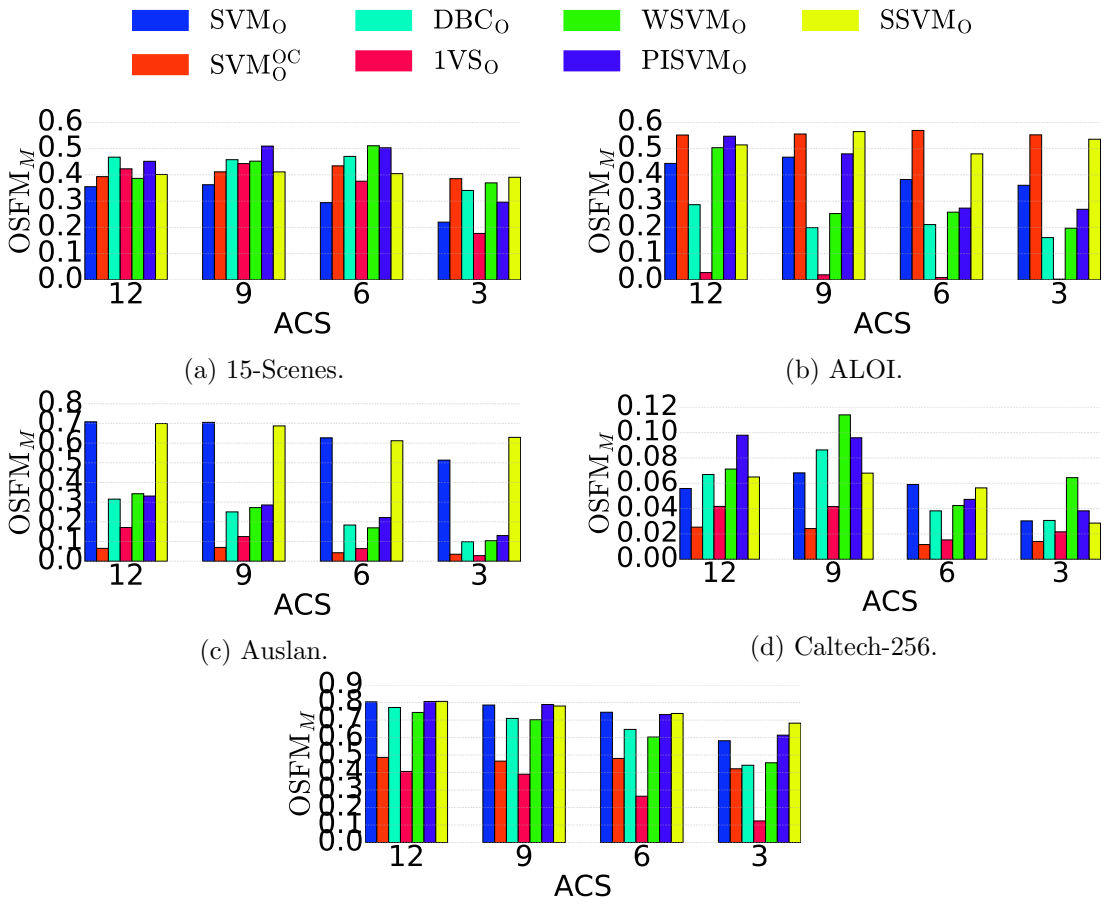

(a) 15-Scenes.

(b) ALOI.

(c) Auslan.

(d) Caltech-256.

(e) Letter.

Figure 6: Results for all data sets regarding macro-averaging open-set f-measure (OSFM$_M$) considering 3, 6, 9, and 12 available classes (ACS).

Measure	SVM	SVM^{OC}	DBC	1VS	WSVM	PISVM
NA	**0.0242**	**0.0000***	**0.0001***	**0.0000***	**0.0003***	0.0973
$OSFM_M$	**0.0385**	**0.0223**	**0.0183**	**0.0003***	**0.0223**	0.1429
$OSFM_\mu$	**0.0249**	0.0655	**0.0013***	**0.0000***	**0.0001***	0.0655
FM_M	0.0655	**0.0183**	**0.0183**	**0.0002***	**0.0321**	0.2024
FM_μ	**0.0034***	**0.0034***	**0.0000***	**0.0000***	**0.0000***	**0.0001***

Table 3: Wilcoxon statistical tests comparing the $SSVM_O$ with other methods using the open-set grid search. Each cell in the table compares the results for all data sets considering all number of available classes. For each number of available classes and data set, the mean of the 10 experiments was taken before the statistical test. **Bold** means there is statistical difference with 95% of confidence. The "*" indicates the statistical difference is with 99% of confidence.

The binomial test takes into account how often the proposed SSVM obtains better results compared to a baseline or, in other words, how often it wins when compared to a given baseline. We also performed the Wilcoxon statistical test to verify if the final classification accuracy obtained by the SSVM is also statistically significant better compared to the baselines or, in other words, to check if the difference in accuracy is significant. For the Wilcoxon test, we used the mean of the 10 experiments for each pair of data set and method and we also used the Holm method for p-value correction. In Table 3, we present the comparison among the methods using the open-set grid search. Once again, in virtually all cases, the $SSVM_O$ outperformed (with statistical significance) all the baselines, except when compared to $PISVM_O$ for which there is statistical difference only regarding FM_μ. For the comparison among the methods using a closed-set grid search (not shown), there are statistical difference in all cases, including for $PISVM_C$. It is worth noticing that Jain et al. (2014) proposed PISVM with the cross-class validation scheme, which has the same principle of the open-set grid search presented herein.

Finally, we also investigated the efficacy of the open-set grid search by comparing the open- and closed-set grid search version of each classifier. We present this comparison in Table 4 for the Binomial test. In virtually all cases, the versions with the open-set grid search obtained statistically significantly better results. Only for the SVM classifier, regarding NA, the version with closed-set grid search obtained better results. However, for this case the difference was not statistically significant.[8]

4.5 Remarks

Note that even the traditional SVM obtained better results than some of the methods in the literature specially tailored for the open-set scenario. We observed that the higher the γ parameter, the better the SVM's performance for the open-set scenario. When grid-searching the γ parameter, the SVM implementation used in this work chooses the higher

8. The raw results obtained in our experiments as well as the script to perform the statistical analysis are available at `https://github.com/pedrormjunior/ssvm-results`.

Measure	SVM	SVM^{OC}	DBC	1VS	WSVM	PISVM	SSVM
NA	*0.8238*	**0.0026***	**0.0004***	1.0000	**0.0000***	**0.0000***	0.8238
$OSFM_M$	0.1153	0.5034	**0.0118**	0.1153	**0.0004***	**0.0026***	**0.0414**
$OSFM_\mu$	0.2632	**0.0118**	**0.0026***	1.0000	**0.0004***	**0.0000***	**0.0118**
FM_M	0.2632	0.1153	0.5034	**0.0118**	**0.0414**	**0.0000***	0.2632
FM_μ	**0.0004***	**0.0118**	**0.0000***	**0.0414**	**0.0000***	**0.0000***	**0.0004***

Table 4: Binomial statistical tests for the pairwise comparison between open- and closed-set grid search implementation of the methods. Each cell in table compares the results for all data sets considering all number of available classes. **Bold** means there is statistical difference with 95% of confidence. The "*" indicates the statistical difference is with 99% of confidence. *Emphasized* means the version with closed-set grid search obtains better performance (only one case), but the difference is not statistically significant.

γ value in case of tie in the validation set. In many cases, there is statistical difference between this implementation and the implementation that chooses the smallest γ in case of ties. Therefore, the SVM used herein is an optimized version for the open-set scenario and it is expected to outperform the traditional SVM implementation.

Another remark is about the frequency of negative bias terms in the binary classifiers that compose the multiclass-from-binary one-vs-all SVM. Most of the binary classifiers for the one-vs-all approach already have the correct negative bias term, as shown in Table 5. To better understand the reason, we also obtained the frequency of binary classifiers with negative bias term using the one-vs-one approach, also shown in Table 5. In this case, only about half of the binary classifiers have negative bias term.

An informal explanation for this behavior is that in the one-vs-all approach, we have more negative than positive samples (and more than one class in the negative set). Then, it is more likely to have the negative samples "around" the positive ones helping the SVM to create a bounded PLOS for the positive class. This intuition is confirmed by Figure 1. For both class 1 and class 2, the SVM creates a bounded PLOS (negative bias term) because class 3 is negative for those binary classifiers and is "around" the positive class in both cases. Considering class 3 as positive, we have no negative samples around the positive class. That is why, in this case, the PLOS is unbounded (non-negative bias term).

The high frequency of negative bias term for the one-vs-all explains why some authors in the literature have been reporting good accuracy for detection problems using SVMs with RBF kernels. For a detection problem, we have one class of interest and multiple others that we consider as negative for what we have access to train with. As the number of negative samples is usually larger, it is more likely to have a classifier with bounded PLOS for detection problems.

Despite most of the cases the SVM obtains the correct bias term for the one-vs-all approach, the optimization problem presented in Section 3 also optimizes the risk of the unknown. That is, it optimizes for the classification in open-set scenario.

Approach	SVM one-vs-all	SVM one-vs-one
15-Scenes	99.00%	56.92%
ALOI	97.67%	52.50%
Auslan	99.00%	42.50%
Caltech-256	98.00%	48.75%
Letter	99.67%	55.75%

Table 5: Percentage of binary classifiers with negative bias term for one-vs-all and one-vs-one approaches.

5. Conclusions and future work

In this paper, we gave sufficient and necessary condition for the the SVM with RBF kernel to have finite risk of the unknown. We then showed that by reformulating the RBF SVM optimization policy to simultaneously optimize margin and ensure a negative bias term, the risk of the unknown is bounded and it provides a formal open-set recognition algorithm.

The proposed SSVM extends upon the traditional SVM's optimization problem. The objective function is changed in the dual problem, but the Lagrangian for the dual problem remains the same. In the dual problem for the SSVM, only a single constraint differs from the SVM's dual problem. The same SMO (Platt, 1998) algorithm can be used to ensure the new constraint is satisfied between iterations. Also, the same WSS (Fan et al., 2005) algorithm can be applied. Consequently, it makes the method easy to extend from an SVM implementation.

As future work, we will investigate alternative forms for ensuring a bounded PLOS for specific implementations of SVMs that do not rely on the bias term, as the ones in Vogt (2002) and Kecman et al. (2005). As shown in Section 3.2, the SVM without the bias term cannot ensure a bounded PLOS, as it depends on the shape of the training data. However, a simple solution can be obtained by training the SVM without bias term and establishing an artificial negative bias term in the decision function. In this case, research can be done on how to obtain this artificial bias term.

In another future work, we will investigate the properties of the multiclass-from-binary SVM with the one-vs-one approach. In a binary SVM, at least one of the two classes will always have infinite PLOS (if one is bounded the other must include all the remaining space and so it must be infinite). We observed experimentally that, according to the way the probability is calculated for each binary classifier (Platt, 2000; Lin et al., 2007) and according to the way the probability estimates are combined in the multiclass-from-binary level (Wu et al., 2004), depending on the threshold established, a bounded KLOS can occur but cannot be ensured. The future work relies on investigating mechanisms to always ensure a bounded KLOS and, consequently, a limited risk of the unknown. This is worth investigating because some works in the literature have presented better results with the one-vs-one than one-vs-all approach for closed-set problems (Galar et al., 2011). We already performed some initial experiments using the SSVM proposed herein as a binary classifier for the one-vs-one approach (rather than the one-vs-all's configuration) and we checked that by ensuring a

bounded PLOS for every binary classifier does not ensure a bounded KLOS in this new configuration. We then hypothesize that, for one-vs-one approach, this investigation should be accomplished in the probability estimation or/and probability combination.

Acknowledgments

The authors thank the financial support of the Brazilian National Council for Scientific and Technological Development (CNPq), the Brazilian Coordination for the Improvement of Higher Level Education Personnel (CAPES) through DeepEyes project and the scholarship provided to the first author, the São Paulo Research Foundation (FAPESP) through DéjàVu project, Grant #2015/19222-9, and Microsoft Research. This research is based upon work funded in part by NSF IIS-1320956.

Appendix A. Complete SSVM formulation

The optimization problem for the SSVM classifier is defined as

$$\begin{aligned} \min_{\mathbf{w},b,\xi} \ & \frac{1}{2}||\mathbf{w}||^2 + C\sum_{i=1}^{m}\xi_i + \lambda b, \\ \text{s.t. } & y_i\left(\mathbf{w}^T\mathbf{x}_i + b\right) - 1 + \xi_i \geq 0, \\ & \xi_i \geq 0. \end{aligned}$$

Using the Lagrangian method, we have the Lagrangian defined as

$$\begin{aligned} \mathcal{L}(\mathbf{w}, b, \xi, \alpha, r) = \frac{1}{2}||w||^2 + C\sum_{i=1}^{m}\xi_i + \lambda b - \sum_{i=1}^{m} r_i\xi_i \\ - \sum_{i=1}^{m}\alpha_i\left[y_i\left(\mathbf{w}^T\mathbf{x}_i + b\right) - 1 + \xi_i\right], \end{aligned}$$

in which $\alpha_i \in \mathbb{R}$ and $r_i \in \mathbb{R}$, $i = 1, \ldots, m$, are the Lagrangian multipliers

First we want to minimize with respect to $\mathbf{w}$, b, and ξ_i, then we must ensure

$$\nabla_{\mathbf{w}} = \frac{\partial}{\partial b}\mathcal{L} = \frac{\partial}{\partial \xi_i}\mathcal{L} = 0. \tag{14}$$

Consequently, we have

$$w - \sum_{i=1}^{m}\alpha_i y_i \mathbf{x}_i = 0 \implies w = \sum_{i=1}^{m}\alpha_i y_i \mathbf{x}_i, \tag{15}$$

$$\lambda - \sum_{i=1}^{m}\alpha_i y_i = 0 \implies \sum_{i=1}^{m}\alpha_i y_i = \lambda, \tag{16}$$

$$C - \alpha_i - r_i = 0 \implies r_i = C - \alpha_i. \tag{17}$$

As the Lagrangian multipliers α_i, r_i must be greater than 0, from Equation (17) we have the constraint $0 \leq \alpha_i \leq C$ as a consequence, in the dual problem, of the soft margin

formulation. This is the same constraint we have in the traditional formulation of the SVM classifier.

Using Equations (15)–(17) to simplify the Lagrangian in Equation (14), we have

$$\mathcal{L}(\mathbf{w}, b, \xi, \alpha, r) = \sum_{i=1}^{m} \alpha_i - \frac{1}{2}||\mathbf{w}||^2,$$

i.e., the same Lagrangian of the traditional SVM optimization problem. The optimization of the bias term b relies on the constraint in Equation (16).

Therefore, the dual optimization problem is defined as

$$\begin{aligned} &\min_{\alpha} W(\alpha) = -\mathcal{L}(\mathbf{w}, b, \xi, \alpha, r) = \frac{1}{2}||\mathbf{w}||^2 - \sum_{i=1}^{m} \alpha_i, \\ &\text{s.t. } 0 < \alpha_i < C, \ \forall i, \\ &\quad \sum_{i=1}^{m} \alpha_i y_i = \lambda. \end{aligned}$$

Appendix B. Additional examples, results and discussion

In this section, we present additional results to the ones presented in Sections 4.3 and 4.4. In Figures 7–9, we present more examples of decision region of different classifiers for 2-d data sets, complementary to those presented in Section 4.3. In Section 4.4, we presented the results for NA and OSFM_M. In Figures 10–12, we present complementary results for different methods with open-set grid search regarding OSFM_μ, FM_M, and FM_μ, respectively.

Depending on the shape of the data set, the SVM is not able to bound the PLOS, as in Figure 7a for the blue class, while for other data sets, e.g., in Figures 8a and 9a, we observe that SVM is able to bound the PLOS for each know class. In Figures 7b–9b, we can observe the highly-specialized behavior of the SVM^{OC}. Furthermore, as we can see in Figure 7b, the SVM^{OC} does not have a good separation among the classes, as it does not take into account the information of the negative classes for each classifier that composes the one-vs-all approach. In Figures 7c and 8c, we notice that the behavior of the DBC is similar to the SVM's one, as it simply moves the hyperplane towards the positive class. As the translation of the hyperplane is based on the empirical risk, measured on the training data for each binary classifier, we can see in Figure 9c that some binary classifiers dominate over others when combining them for the multiclass-from-binary extension. As expected, see Figures 7d–9d, 1VS leaves an unbounded KLOS. However, the confusion showed in Figure 9d is because the data set is only 2-dimensional. In Figures 7e–9e, we can notice the prevalence of the one-class model of the WSVM, however not as specialized as the SVM^{OC}, due to the parameters adjustment performed by WSVM. We also note in Figures 7e and 8e that the binary model of the WSVM is able to better separate among the known classes, compared to SVM^{OC} in Figure 7b. In Figures 8f and 8f, we can see that PISVM is able the bound the PLOS for every class, however, as seen in Figure 7f, that is not always the case. The authors of PISVM proposed the method to optimize according to the risk of the unknown, however they did not prove a bounded KLOS for the method. Finally, in Figures

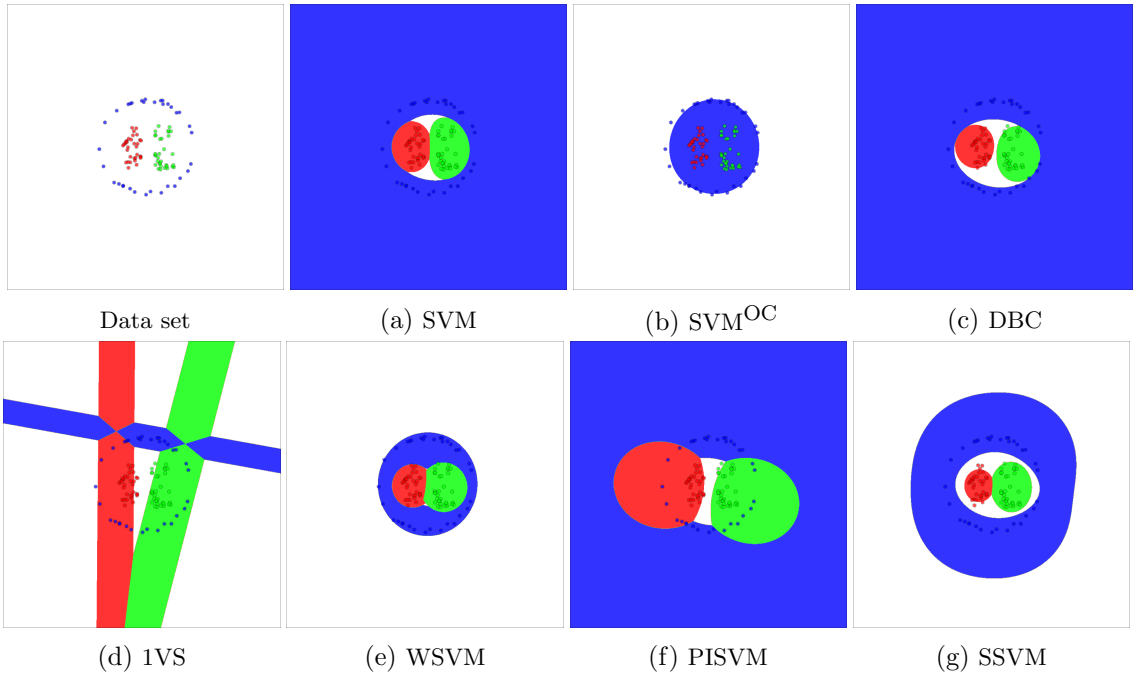

Data set (a) SVM (b) SVM^{OC} (c) DBC

(d) 1VS (e) WSVM (f) PISVM (g) SSVM

Figure 7: Decision boundaries for the Boat data set. The non-white regions represent the region in which a test sample would be classified as belonging to the same class of the samples with the same color. All samples in the white regions would be classified as unknown.

7g–9g, we observe the behavior of the proposed SSVM and how it gracefully bounds the PLOS.

In Figures 10–12, for the results regarding OSFM$_\mu$, FM$_M$, and FM$_\mu$, we see that SSVM$_\text{O}$ does not obtain the best behavior in all cases but it consistently appears in the top positions for all the data sets, while the other methods have better results for one or two data sets and lower performance for the others. Overall, the SSVM$_\text{O}$ has better performance with statistical difference for those measures, as shown in Tables 2 and 3.

References

Abhijit Bendale and Terrance E. Boult. Towards open world recognition. In *Intl. Conference on Computer Vision and Pattern Recognition*, pages 1893–1902, Boston, MA, June 2015. IEEE Press.

Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Information Science and Statistics. Springer-Verlag New York, 1 edition, 2006.

Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. A training algorithm for optimal margin classifiers. In *Annual Workshop on Computational Learning Theory*, pages 144–152, Pittsburgh, PA, July 1992. ACM Press.

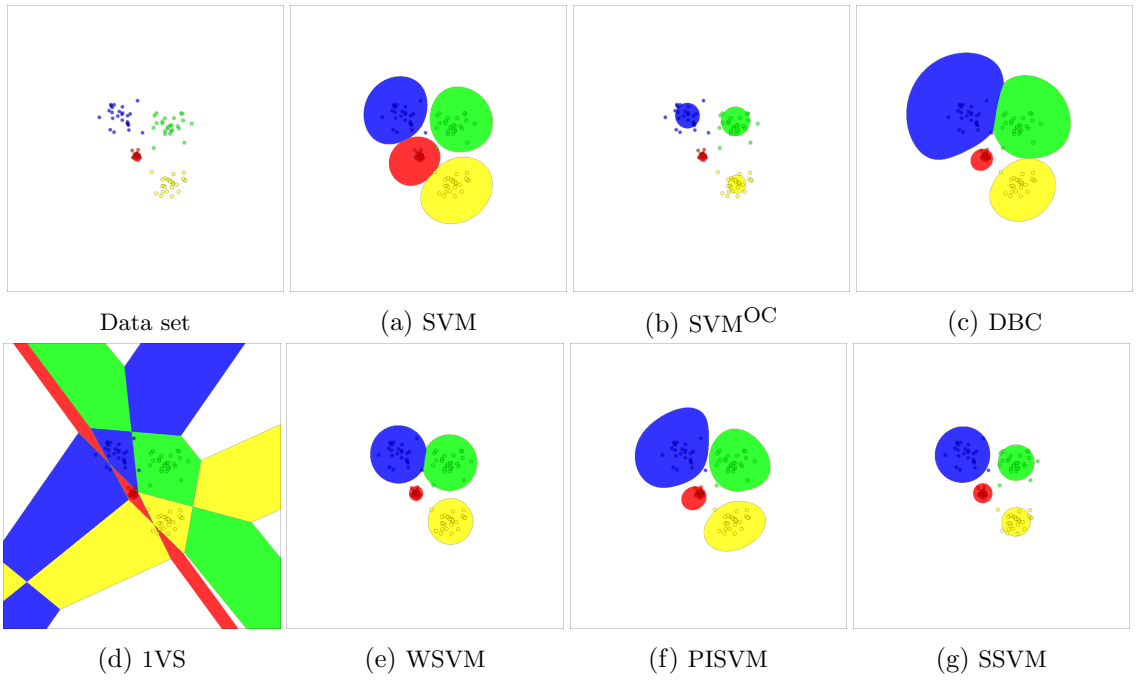

Data set (a) SVM (b) SVM^{OC} (c) DBC

(d) 1VS (e) WSVM (f) PISVM (g) SSVM

Figure 8: Decision boundaries for the Four-gauss data set. The non-white regions represent the region in which a test sample would be classified as belonging to the same class of the samples with the same color. All samples in the white regions would be classified as unknown.

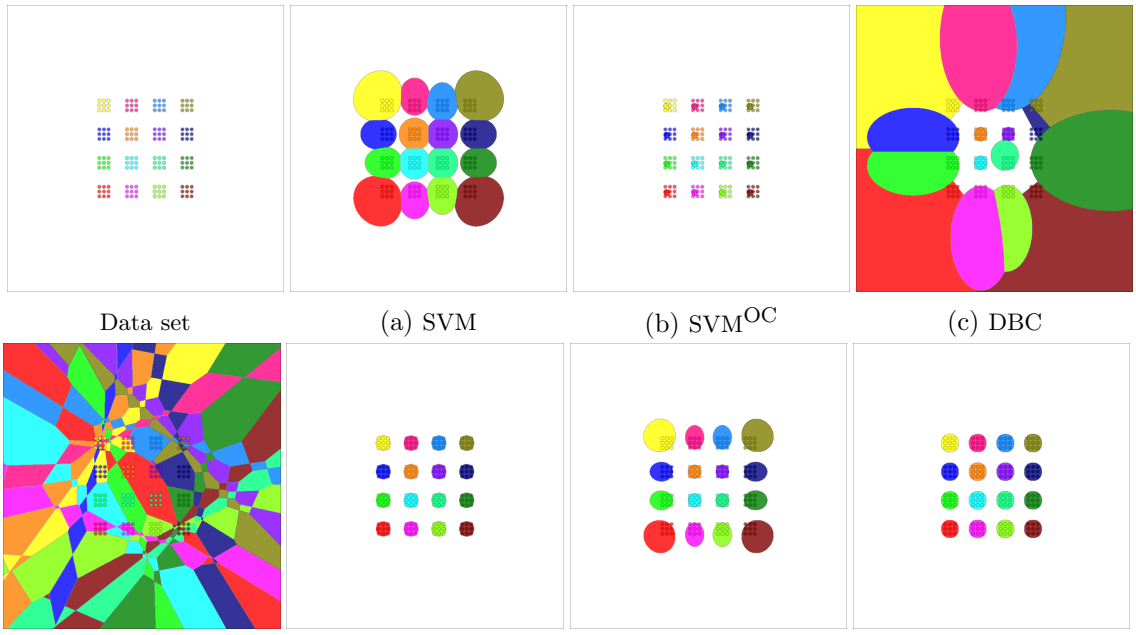

Data set (a) SVM (b) SVM^{OC} (c) DBC

(d) 1VS (e) WSVM (f) PISVM (g) SSVM

Figure 9: Decision boundaries for the Regular data set. The non-white regions represent the region in which a test sample would be classified as belonging to the same class of the samples with the same color. All samples in the white regions would be classified as unknown.

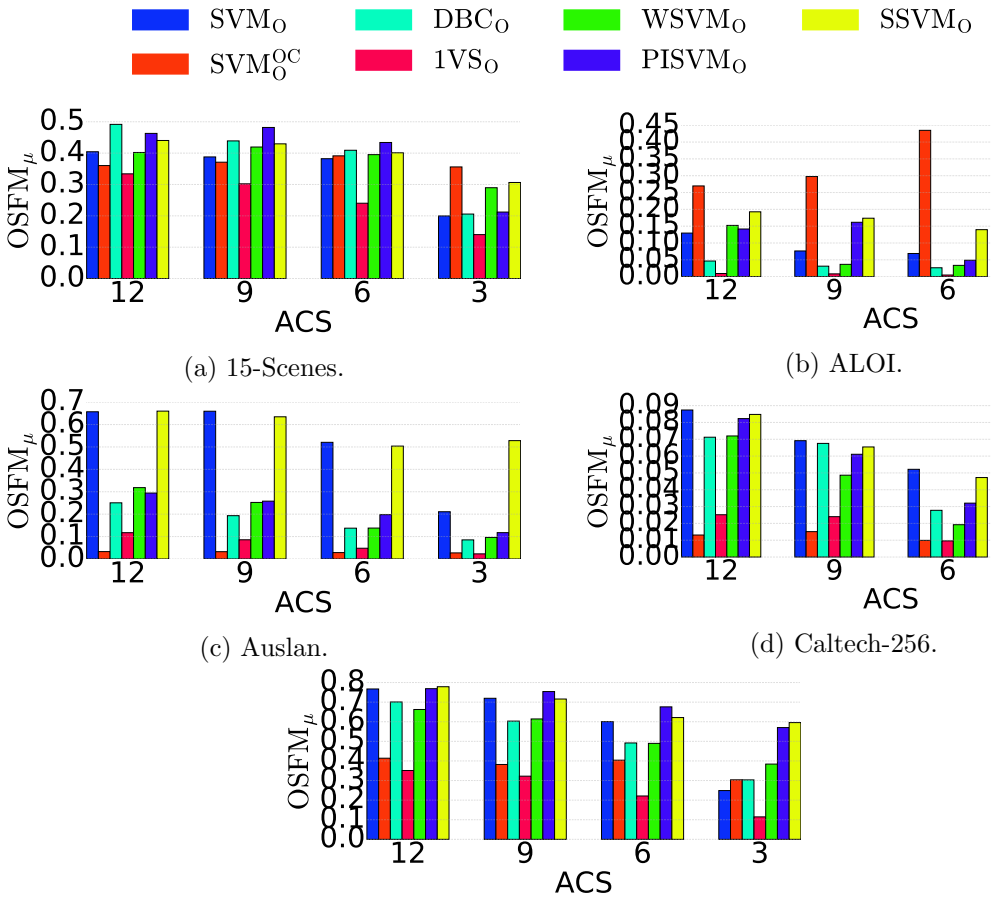

(a) 15-Scenes.

(b) ALOI.

(c) Auslan.

(d) Caltech-256.

(e) Letter.

Figure 10: Results for all data sets regarding micro-averaging open-set f-measure (OSFM_μ) considering 3, 6, 9, and 12 available classes (ACS).

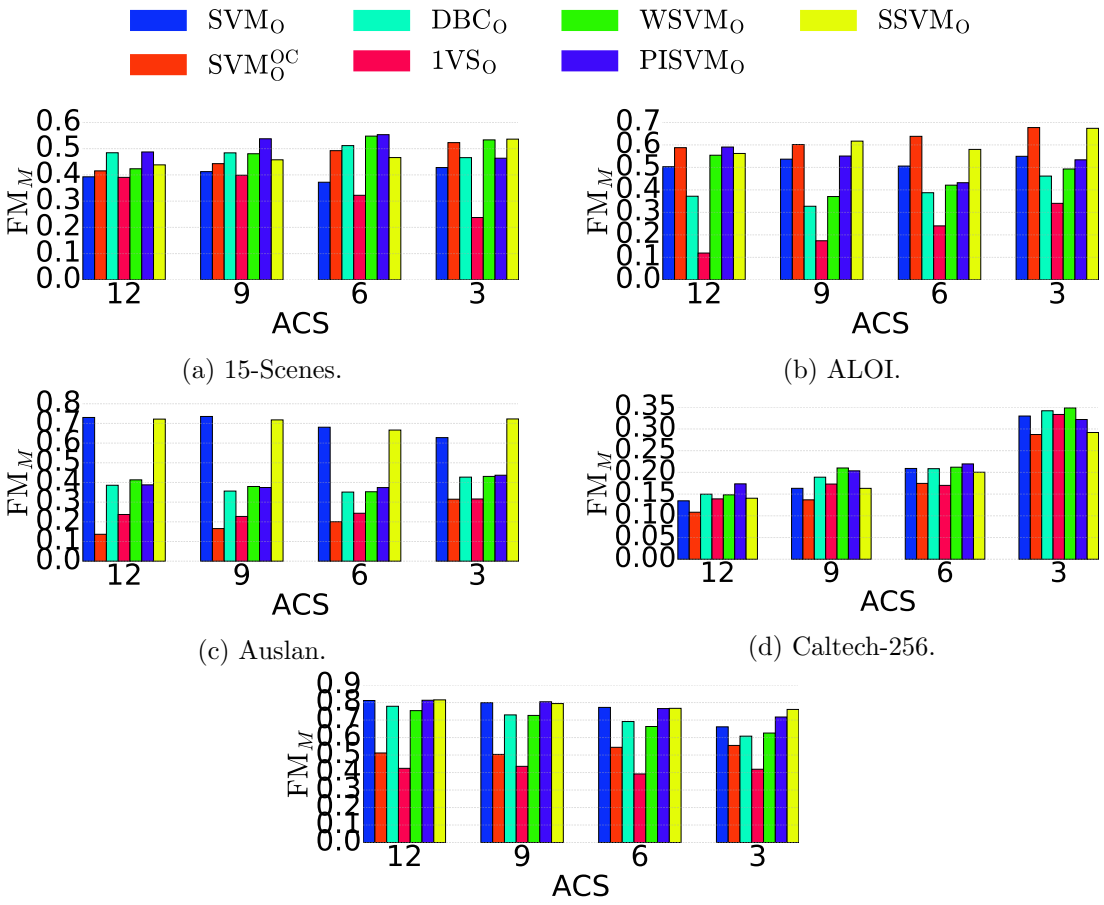

(a) 15-Scenes.

(b) ALOI.

(c) Auslan.

(d) Caltech-256.

(e) Letter.

Figure 11: Results for all data sets regarding macro-averaging f-measure (FM_M) considering 3, 6, 9, and 12 available classes (ACS).

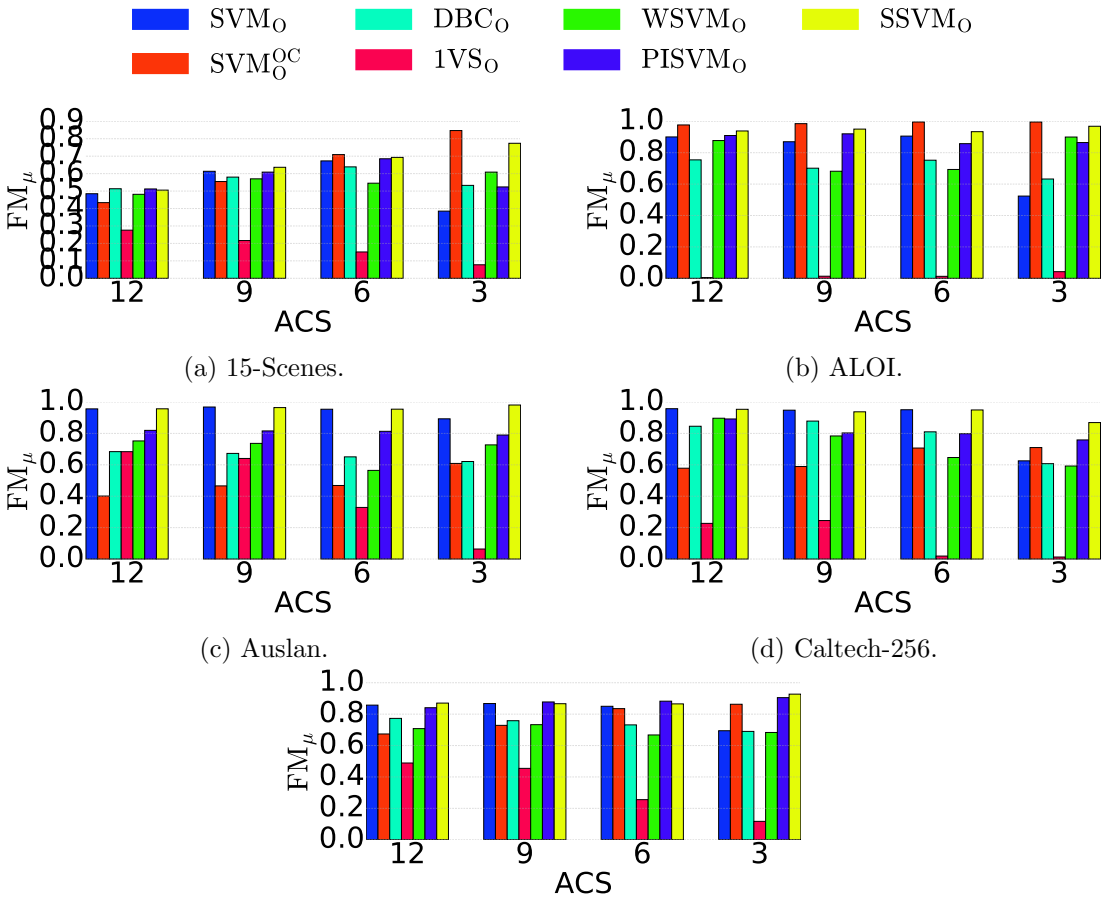

(a) 15-Scenes.

(b) ALOI.

(c) Auslan.

(d) Caltech-256.

(e) Letter.

Figure 12: Results for all data sets regarding micro-averaging f-measure (FM_μ) considering 3, 6, 9, and 12 available classes (ACS).

Y-Lan Boureau, Francis Bach, Yann LeCun, and Jean Ponce. Learning mid-level features for recognition. In *Intl. Conference on Computer Vision and Pattern Recognition*, pages 2559–2566, San Francisco, CA, June 2010. IEEE Press.

Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, October 2001.

Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *Transactions on Intelligent Systems and Technology*, 2(3):27:1–27:27, April 2011.

Stuart Coles. *An introduction to statistical modeling of extreme values*. Springer Series in Statistics. Springer, 1 edition, 2001.

Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3): 273–297, September 1995.

Filipe de Oliveira Costa, Michael Eckmann, Walter J. Scheirer, and Anderson Rocha. Open set source camera attribution. In *Conference on Graphics, Patterns, and Images*, pages 71–78, Ouro Preto, August 2012. IEEE Press.

Filipe de Oliveira Costa, Ewerton Silva, Michael Eckmann, Walter J. Scheirer, and Anderson Rocha. Open set source camera attribution and device linking. *Pattern Recognition Letters*, 39:92–101, April 2014.

Janez Demšar. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, 7:1–30, December 2006.

Rong-En Fan, Pai-Hsuen Chen, and Chih-Jen Lin. Working set selection using second order information for training support vector machines. *Journal of Machine Learning Research*, 6(dec):1889–1918, December 2005.

Peter W. Frey and David J. Slate. Letter recognition using holland-style adaptive classifiers. *Machine Learning*, 6(2):161–182, March 1991.

Mikel Galar, Alberto Fernndez, Edurne Barrenechea, Humberto Bustince, and Francisco Herrera. An overview of ensemble methods for binary classifiers in multi-class problems: Experimental study on one-vs-one and one-vs-all schemes. *Pattern Recognition*, 44(8): 1761–1776, August 2011.

Jan-Mark Geusebroek, Gertjan J. Burghouts, and Arnold W. M. Smeulders. The amsterdam library of object images. *Intl. Journal of Computer Vision*, 61(1):103–112, January 2005.

Greg Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. Technical report, California Institute of Technology, 2007.

Brian Heflin, Walter J. Scheirer, and Terrance E. Boult. Detecting and classifying scars, marks, and tattoos found in the wild. In *Intl. Conference on Biometrics: Theory, Applications and Systems*, pages 31–38, Arlington, VA, September 2012. IEEE Press.

Gordon F. Hughes. On the mean accuracy of statistical pattern recognizers. *Transactions on Information Theory*, 14(1):55–63, January 1968.

Lalit P. Jain, Walter J. Scheirer, and Terrance E. Boult. Multi-class open set recognition using probability of inclusion. In *European Conference on Computer Vision*, volume 8691 of *Lecture Notes in Computer Science*, pages 393–409. Springer International Publishing, Zurich, Switzerland, September 2014.

Mohammed Waleed Kadous. *Temporal classification: Extending the classification paradigm to multivariate time series*. PhD thesis, The University of New South Wales, New South Wales, Australia, October 2002.

V. Kecman, T.-M. Huang, and M. Vogt. *Support Vector Machines: Theory and Applications*, chapter Iterative single data algorithm for training kernel machines from huge data sets: Theory and performance, pages 255–274. Studies in Fuzziness and Soft Computing. Springer Berlin Heidelberg, April 2005.

Ludniila I. Kuncheva and Stefan T. Hadjitodorov. Using diversity in cluster ensembles. In *Intl. Conference on Systems, Man and Cybernetics*, volume 2, pages 1214–1219, The Hague, Netherlands, October 2004. IEEE Press.

Svetlana Lazebnik, Cordelia Schmid, and Jean Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *Intl. Conference on Computer Vision and Pattern Recognition*, volume 2, pages 2169–2178, New York, NY, June 2006. IEEE Press.

Hsuan-Tien Lin, Chih-Jen Lin, and Ruby C. Weng. A note on platt's probabilistic outputs for support vector machines. *Machine Learning*, 68(3):267–276, October 2007.

David Lowe. Distinctive image features from scale-invariant keypoints. *Intl. Journal of Computer Vision*, 60(2):91–110, November 2004.

D. Michie, D. J. Spiegelhalter, and C. C. Taylor. *Machine learning, neural and statistical classification*. Ellis Horwood Series in Artificial Intelligence. Prentice Hall, Upper Saddle River, NJ, July 1994.

John C. Platt. Fast training of support vector machines using sequential minimal optimization. In Bernhard Schölkopf, Christopher J. C. Burges, and Alexander J. Smola, editors, *Advances in kernel methods: Support vector learning*, pages 185–208. MIT Press, December 1998.

John C. Platt. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. In Alexander J. Smola, Peter Bartlett, Bernhard Schölkopf, and Dale Schuurmans, editors, *Advances in Large Margin Classifiers*, Neural Information Processing, pages 61–74. MIT Press, September 2000.

Dimitrios A. Pritsos and Efstathios Stamatatos. Open-set classification for automated genre identification. In Pavel Serdyukov, Pavel Braslavski, Sergei O. Kuznetsov, Jaap Kamps, Stefan Rüger, Eugene Agichtein, Ilya Segalovich, and Emine Yilmaz, editors, *Advances in Information Retrieval*, volume 7814 of *Lecture Notes in Computer Science*, pages 207–217. Springer Berlin Heidelberg, March 2013.

Anderson Rocha and Siome Goldenstein. Multiclass from binary: Expanding one-vs-all, one-vs-one and ECOC-based approaches. *Transactions on Neural Networks and Learning Systems*, 25(2):289–302, February 2014.

Walter J. Scheirer, Anderson Rocha, Ross J. Micheals, and Terrance E. Boult. Meta-recognition: The theory and practice of recognition score analysis. *Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1689–1695, August 2011.

Walter J. Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E. Boult. Towards open set recognition. *Transactions on Pattern Analysis and Machine Intelligence*, 35(7):1757–1772, July 2013.

Walter J. Scheirer, Lalit P. Jain, and Terrance E. Boult. Probability models for open set recognition. *Transactions on Pattern Analysis and Machine Intelligence*, 36(11):2317–2324, November 2014.

Bernhard Schölkopf and Alex J. Smola. *Learning with kernels*. Adaptive Computation and Machine Learning. The MIT Press, 1 edition, December 2001.

Bernhard Schölkopf, John C. Platt, John Shawe-Taylor, Alex J. Smola, and Robert C. Williamson. Estimating the support of a high-dimensional distribution. Technical report, Microsoft Research, Redmond, WA, November 1999.

Marina Sokolova and Guy Lapalme. A systematic analysis of performance measures for classification tasks. *Information Processing and Management*, 45(4):427–437, July 2009.

Renato O. Stehling, Mario A. Nascimento, and Alexandre Xavier Falcão. A compact and efficient image retrieval approach based on border/interior pixel classification. In *Intl. Conference on Information and Knowledge Management*, pages 102–109, McLean, VA, November 2002. ACM Press.

Jan C. van Gemert, Cor J. Veenman, Arnold W. M. Smeulders, and Jan-Mark Geusebroek. Visual word ambiguity. *Transactions on Pattern Analysis and Machine Intelligence*, 32 (7):1271–1283, July 2010.

Vladimir Vapnik and Alexey Chervonenkis. A note on one class of perceptrons. *Automation and Remote Control*, 25(1), 1964.

Michael Vogt. SMO algorithms for support vector machines without bias term. Technical report, Technische Universität Darmstadt, July 2002.

Ting-Fan Wu, Chih-Jen Lin, and Ruby C. Weng. Probability estimates for multi-class classification by pairwise coupling. *Journal of Machine Learning Research*, 5:975–1005, December 2004.